

THE STRUCTURE OF THE CONSECUTIVE PATTERN POSET

SERGI ELIZALDE AND PETER R. W. MCNAMARA

ABSTRACT. The consecutive pattern poset is the infinite partially ordered set of all permutations where $\sigma \leq \tau$ if τ has a subsequence of adjacent entries in the same relative order as the entries of σ . We study the structure of the intervals in this poset from topological, poset-theoretic, and enumerative perspectives. In particular, we prove that all intervals are rank-unimodal and strongly Sperner, and we characterize disconnected and shellable intervals. We also show that most intervals are not shellable and have Möbius function equal to zero.

CONTENTS

1. Introduction	2
2. Preliminaries	4
2.1. Fundamentals	4
2.2. The Möbius function papers	4
2.3. Initial observations	6
3. Disconnectivity	6
3.1. Characterization of disconnected intervals	7
3.2. Finding disconnected intervals	8
4. Shellability	11
4.1. Background on shellability and CL-shellability	11
4.2. Shellability in the consecutive pattern poset	12
5. All intervals are rank-unimodal and strongly Sperner	16
5.1. Rank-unimodality	16
5.2. Strongly Sperner property	19
6. The exterior of a permutation	21
6.1. The length of the exterior	21
6.2. Asymptotic behavior of the length of the exterior	22
6.3. Permutations with no carrier, and the Möbius function of most intervals	25
Acknowledgement	27
References	27

2010 *Mathematics Subject Classification.* Primary 06A07; Secondary 05A05, 05A16, 05E45, 52B22, 60C05.

Key words and phrases. Consecutive pattern, poset, disconnected, shellable, rank unimodal, Sperner, exterior, Möbius function.

Sergi Elizalde was partially supported by grant #280575 from the Simons Foundation, by grant H98230-14-1-0125 from the NSA, and by the Distinguished Visiting Professor Program at Bucknell University. Peter McNamara was partially supported by grant #245597 from the Simons Foundation.

1. INTRODUCTION

Consecutive patterns in permutations generalize well-studied notions such as descents, ascents, peaks, valleys and runs. A permutation σ is said to be *contained* in another one τ as a *consecutive pattern* if τ has a subsequence of adjacent entries in the same relative order as the entries of σ . Otherwise, τ is said to *avoid* σ . For example, permutations that avoid both 123 and 321 as a consecutive pattern are the well-known alternating and reverse-alternating permutations. A systematic study of enumerative aspects of consecutive patterns started in [EN03] and, in the last decade, such patterns have become a vibrant area of research; see [Eli16] for a recent survey. Underlying all these questions is a partial order \mathcal{P} on the set of all permutations, where we define $\sigma \leq \tau$ if σ is contained in τ as a consecutive pattern. This paper is the first systematic study of intervals in this consecutive pattern poset. See Figure 1.1 for an example of such an interval.

Some questions for consecutive patterns are motivated by the analogous problems for so-called classical patterns, one of the most actively studied topics of combinatorics in the last three decades. For the definition of containment in this classical case, we remove the restriction that the relevant subsequence of τ consists of adjacent entries. For example, 123 is less than 2314 in the classical pattern poset but not in the consecutive pattern poset. See [Bón12] for an exposition of some of the most important developments in the study of classical patterns, [Kit11] for a detailed treatment, and [Ste13] for a survey of recent developments.

Examples of questions for consecutive patterns that are motivated by the classical case include classifying them into equivalence classes determined by the number of permutations avoiding them [Nak11], finding generating functions for the distribution of occurrences of a fixed pattern in permutations [EN03, MR06, EN12], and determining the asymptotic growth of the number of permutations avoiding a pattern [Eli13, Per13]. Consecutive patterns are interesting not only because the answers to the above questions and the techniques used to solve them are usually quite different than for classical patterns, but also because they have important applications to dynamical systems [AEK08, Eli09].

While most of the aforementioned work is enumerative, our approach also has a poset-theoretic and topological flavor. An early inspiration for research from this viewpoint was a question of Wilf [Wil02] asking for the Möbius function of intervals in the poset defined by classical pattern containment. This question remains wide open, but has received increasing attention of late [BJJS11, MS15, SV06, Smi14, Smi15, Smi16, ST10]. In contrast, the Möbius function for intervals in \mathcal{P} , the consecutive pattern poset, has been determined by Bernini, Ferrari and Steinrímsson [BFS11] and by Sagan and Willenbring [SW12]. This already gives an indication that the consecutive pattern case is more tractable for certain types of questions than the classical case.

The precursor in the classical case to the present work is [MS15], where the focus is on classifying disconnected open intervals and showing that certain special intervals are shellable. We successfully address these same topological questions of disconnectivity and shellability for the consecutive pattern poset \mathcal{P} . Furthermore, we consider what poset-theoretic properties might hold for \mathcal{P} , e.g., we show that all intervals are rank-unimodal, a statement that is just a conjecture in the classical case [MS15]. Recall that a finite graded poset of rank N with a_i elements of rank i is called *rank-unimodal* if the sequence a_0, a_1, \dots, a_N is unimodal, meaning that

there exists k with $0 \leq k \leq N$ such that $a_0 \leq a_1 \leq \cdots \leq a_k \geq a_{k+1} \geq \cdots \geq a_N$. Our main results include the following.

- A simple classification of those open intervals in \mathcal{P} that are disconnected (Theorem 3.2).
- Almost all intervals contain a disconnected subinterval of rank at least 3, in a certain precise sense, and are thus not shellable (Theorem 3.6 and Corollary 4.2).
- All other intervals are shellable (Theorem 4.3). One motivation for shellability is that it completely determines the homotopy type of the order complex of the interval as a wedge of spheres. The number of spheres is the absolute value of the Möbius function and is thus readily determined using [BFS11, SW12].
- All intervals are rank-unimodal (Corollary 5.2). The highest rank of the maximum cardinality is easily determined and the interval takes on a particularly simple structure above this rank.
- All intervals are strongly Sperner (Theorem 5.6), a condition equating the sizes of the union of the k largest ranks and the largest union of k antichains, for all k .
- It is clear from the formula for the Möbius function $\mu(\sigma, \tau)$ that it depends heavily on the *exterior* of τ , which is the longest permutation that is both a proper prefix and proper suffix of τ . We initiate a study of the asymptotic behavior of the exterior in Section 6. We show that its expected length is bounded between $e - 1$ and e (Theorem 6.5).
- Almost all intervals have a zero Möbius function, in a certain precise sense (Corollary 6.10).

Taken together, these topological, poset-theoretic and enumerative results give a rather comprehensive picture of the structure of intervals in \mathcal{P} .

Preliminaries and some initial observations are the content of the next section. Disconnectivity is the subject of Section 3, shellability is addressed in Section 4, and the properties of rank-unimodality and strongly Sperner are shown in Section 5. The exterior of a permutation is considered in Section 6, along with several open problems about its behavior.




FIGURE 1.1. The interval $[12, 213546]$ in \mathcal{P} .

2. PRELIMINARIES

In this section, we collect together some useful terminology and notation, and make some initial observations.

2.1. Fundamentals. Let \mathcal{S}_n denote the set of permutations of $[n] := \{1, 2, \dots, n\}$. If $\tau \in \mathcal{S}_n$, we denote the length of τ by $|\tau| = n$, and we write τ in one-line notation as $\tau = \tau_1 \tau_2 \dots \tau_n$. Given a sequence $a_1 a_2 \dots a_k$ of distinct positive integers, define its *reduction* $\rho(a_1 a_2 \dots a_k)$ to be the permutation of $\{1, \dots, k\}$ obtained by replacing the smallest entry with 1, the second smallest with 2, and so on. For example, $\rho(394176) = 263154$. For $1 \leq i \leq j \leq n$, let $\tau_{[i,j]} = \rho(\tau_i \dots \tau_j)$, and note that $\tau_{[i,j]} \in \mathcal{S}_{j-i+1}$.

We can now define our poset of interest. We say that τ *contains* σ as a consecutive pattern if there exist $1 \leq i \leq j \leq n$ such that $\tau_{[i,j]} = \sigma$. We write $\sigma \leq \tau$, and we say that $\tau_i \dots \tau_j$ is an occurrence of σ . The relation \leq defines a partial order on the set $\mathcal{S} = \bigcup_{n \geq 1} \mathcal{S}_n$ of all permutations, and we denote by \mathcal{P} the corresponding partially ordered set, called the *consecutive pattern poset*. The poset \mathcal{P} is the main object of study in this paper. See Figure 1.1 for an example of an interval in \mathcal{P} .

A related partial order on \mathcal{S} is obtained by considering classical pattern containment instead, i.e., the case when τ is said to contain σ if any subsequence of τ (not necessarily in consecutive positions) has reduction equal to σ . We will refer to the resulting poset as the *classical pattern poset*.

In the rest of this paper we will say that τ contains σ to mean that τ contains σ as a consecutive pattern, that is, $\sigma \leq \tau$. If this is not the case, we say that τ avoids σ .

We next recall some common operations on permutations. The *reversal* of the permutation $\tau = \tau_1 \tau_2 \dots \tau_n$ is the permutation $\tau_n \dots \tau_2 \tau_1$. The *complement* of τ is the permutation whose i th entry is $n + 1 - \tau_i$. Note that both the reversal and complement operations preserve the order relation in \mathcal{P} . We note, however, that unlike in the classical pattern poset, the inverse operation does not preserve the order relation in \mathcal{P} ; e.g., the relation $132 \leq 3142$ is not preserved when taking inverses.

Given $\sigma \in \mathcal{S}_m$ and $\tau \in \mathcal{S}_n$, we define their *direct sum* $\sigma \oplus \tau$ as the concatenation of σ and the permutation τ^{+m} formed by adding m to every entry of τ . Similarly, the *skew sum* $\sigma \ominus \tau$ is the concatenation of σ^{+n} and τ . For example, $1324 \oplus 21 = 132465$ and $1324 \ominus 21 = 354621$.

To consider disconnectivity in Section 3 and shellability in Section 4, we need to define a simplicial complex associated with any interval of any poset. Given an interval $[\sigma, \tau]$, its *order complex* $\Delta(\sigma, \tau)$ is the simplicial complex whose faces are the chains of the open interval (σ, τ) . For example, the order complex of the interval $[12, 213546]$ from Figure 1.1 is shown in Figure 2.1. We will postpone the other necessary background about shellability until Section 4. We refer the reader to [Wac07] for background on poset topology.

2.2. The Möbius function papers. It remains a wide open problem to determine the Möbius function for the *classical* pattern poset; see [BJJS11, MS15, SV06, Smi14, Smi15, Smi16, ST10] for special cases. In contrast, the Möbius function for the consecutive pattern poset \mathcal{P} has been determined independently by Bernini, Ferrari and Steingrímsson in [BFS11] and by Sagan and Willenbring in [SW12]. We will present this recursive formula as stated in the latter paper here because it




FIGURE 2.1. The order complex of the interval $[12, 213546]$ in \mathcal{P} from Figure 1.1. Only chains in the *open* interval (σ, τ) are included in the order complex of $[\sigma, \tau]$; including σ and τ would simply result in the join of the order complex with a line segment and hence would be contractible.

introduces some key ideas we will need later, and then we will explain the parts of the viewpoint of the former paper that will be useful to us.

We say that σ is a *prefix* (respectively *suffix*) of $\tau \in \mathcal{S}_n$ if $\sigma = \tau_{[1,i]}$ (resp. $\sigma = \tau_{[n-i+1,n]}$) for some i , and is a *proper* prefix (resp. suffix) if it does not equal τ . For example, 312 is a proper prefix of 51342. We say that σ is a *bifix* of τ if it is both a proper prefix and proper suffix of τ . The *exterior* of τ , denoted $x(\tau)$, is the longest bifix of τ . Note that $x(\tau)$ always exists whenever $\tau > 1$ since 1 is a proper prefix and suffix of τ in that case. The *interior* of τ is $\tau_{[2,n-1]}$ and is denoted $i(\tau)$. For example, if $\tau = 21435$, then $x(\tau) = 213$ and $i(\tau) = 132$. We say that τ is *monotone* if it equals $12\dots n$ or $n\dots 21$.

Theorem 2.1 ([BFS11, SW12]). *For $\sigma, \tau \in \mathcal{P}$ with $\sigma \leq \tau$,*

$$\mu(\sigma, \tau) = \begin{cases} \mu(\sigma, x(\tau)) & \text{if } |\tau| - |\sigma| > 2 \text{ and } \sigma \leq x(\tau) \not\leq i(\tau), \\ 1 & \text{if } |\tau| - |\sigma| = 2, \tau \text{ is not monotone, and } \sigma \in \{i(\tau), x(\tau)\}, \\ (-1)^{|\tau|-|\sigma|} & \text{if } |\tau| - |\sigma| < 2, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear from this important theorem that the exterior of a permutation is worthy of attention, and yet this attention seems to be absent from the literature. This motivates our detailed consideration of the statistic $|x(\tau)|$ for general permutations in Subsections 6.1 and 6.2. In particular, we show that the expected value of $|x(\tau)|$ as $|\tau| \rightarrow \infty$ is bounded between $e - 1$ and e .

In [BFS11], if the exterior $x(\tau)$ satisfies the condition that $\sigma \leq x(\tau) \not\leq i(\tau)$, then $x(\tau)$ is called the *carrier element* of $[\sigma, \tau]$. As we see in Theorem 2.1, when $[\sigma, \tau]$ has a carrier element we get the interesting case where the calculation of $\mu(\sigma, \tau)$ reduces to determining $\mu(\sigma, x(\tau))$. It is therefore natural to ask under what conditions $[\sigma, \tau]$ has a carrier element. This motivates our consideration in Subsection 6.3 of the permutations τ such that $[1, \tau]$ has a carrier element. In this case of $\sigma = 1$, $x(\tau)$ can be called the *carrier element of τ* , and it exists if and only if $x(\tau) \not\leq i(\tau)$.

We determine the asymptotic probability of a permutation not having a carrier element and, as an application, prove that the Möbius function is almost always 0, in a particular precise way.

2.3. Initial observations. One feature of the consecutive pattern poset that makes it more tractable than the classical pattern poset is that any permutation τ of length n covers at most two elements, namely $\tau_{[1,n-1]}$ and $\tau_{[2,n]}$. The next lemma will be useful later; it appears in [BFS11, SW12] and is routine to check.

Lemma 2.2. *For $\tau \in \mathcal{P}$ of length n , we have $\tau_{[1,n-1]} = \tau_{[2,n]}$ if and only if τ is monotone.*

As a consequence, we get our first structural result about intervals in \mathcal{P} .

Proposition 2.3. *The interval $[\sigma, \tau]$ in \mathcal{P} is a chain if and only if either*

- τ is monotone, or
- σ occurs exactly once in τ and does so as a prefix or a suffix.

Proof. The “if” direction is clear, while for the “only if” direction we observe that if $[\sigma, \tau]$ is a chain, then τ covers just one element, leading to the two given possibilities. \square

As previously shown in [BFS11], we can completely determine the structure of $[\sigma, \tau]$ when σ has just one occurrence in τ , a result we prove here for the purposes of self-containment. This contrasts with the classical pattern poset, where the corresponding result is unknown.

Proposition 2.4 ([BFS11]). *For an interval $[\sigma, \tau]$ in \mathcal{P} , if σ occurs exactly once in τ then $[\sigma, \tau]$ is a product of two chains. Moreover, if $\sigma = \tau_{[i,j]}$, then these chains have lengths i and $|\tau| - j + 1$.*

Proof. Any $\pi \in [\sigma, \tau]$ is obtained from τ by deleting a entries from the left end of τ and b entries from the right end of τ , where

$$0 \leq a \leq i-1 \quad \text{and} \quad 0 \leq b \leq |\tau| - j. \quad (2.1)$$

Thus there is a surjective map ϕ from such tuples (a, b) to the elements π of $[\sigma, \tau]$. This map is injective since there is just one occurrence of σ in τ , allowing us to uniquely determine a and b from π . We also see that ϕ is order-reversing in the sense that $\phi(a, b) \leq \phi(a', b')$ in $[\sigma, \tau]$ if and only if $a \geq a'$ and $b \geq b'$ as integers. Therefore, $[\sigma, \tau]$ is order-isomorphic to the dual of a product of two chains, which is itself a product of two chains. The lengths of the chains follow from (2.1). \square

3. DISCONNECTIVITY

We say that an interval $[\sigma, \tau]$ is *disconnected* (resp. *shellable*) if $\Delta(\sigma, \tau)$ is disconnected (resp. shellable). Equivalently, $[\sigma, \tau]$ is disconnected if the Hasse diagram of the open interval (σ, τ) is disconnected. For example, in Figure 1.1, the subinterval $[213, 213546]$ is disconnected. In examining the structure of intervals of \mathcal{P} , a natural question is to ask when such intervals $[\sigma, \tau]$ are disconnected, preferably giving the answer in terms of simple conditions on σ and τ . We answer this question in Theorem 3.2 below. Determining when an interval contains a non-trivial disconnected subinterval is a more difficult task, which we explore following Theorem 3.2. Interestingly, as we show in Theorem 3.6, almost all intervals in \mathcal{P} do contain such a disconnected subinterval.

3.1. Characterization of disconnected intervals. It will be helpful to deal with posets of rank 2 separately because we can completely classify them, and because they sometimes require special treatment. In particular, unlike disconnected intervals of rank at least 3, a disconnected interval of rank 2 is shellable; we will study the topic of shellability in detail in the next section. Since every element of \mathcal{P} covers at most 2 elements, a rank-2 interval in \mathcal{P} is either a chain or has two elements of rank 1. Thus the structure of rank-2 intervals is completely determined by Proposition 2.3. With rank-2 disconnected intervals now classified, we will follow [MS15] in saying that a disconnected interval is *non-trivial* if it has rank at least 3.

Definition 3.1. For $\sigma < \tau$, we say that σ straddles τ if σ is both a prefix and a suffix of τ and has no other occurrences in τ .

It is easy to check that σ straddles τ if and only if σ is the carrier element of $[\sigma, \tau]$, as defined in the last paragraph of Subsection 2.2. Either way we state it, this is exactly the condition that causes non-trivial disconnectedness of $[\sigma, \tau]$, as we now show.

Theorem 3.2. For $\sigma, \tau \in \mathcal{P}$ with $|\tau| - |\sigma| \geq 3$, we have that $[\sigma, \tau]$ is disconnected if and only if σ straddles τ .

Proof. Let $n = |\tau|$. If σ straddles τ , then the open interval (σ, τ) consists of two chains: one containing the elements $\tau_{[1,j]}$ for $|\sigma| < j < n$, and one containing elements $\tau_{[n+1-j,n]}$ for $|\sigma| < j < n$. These chains are disjoint since σ appears as a prefix in the elements of the former chain, and as a suffix in the elements of the latter chain. Thus (σ, τ) is disconnected, as required. This argument is very similar to that of [BFS11, Lemma 3.3], where their conclusion is that $\mu(\sigma, \tau) = 1$.

To prove the converse, suppose now that (σ, τ) is disconnected. In \mathcal{P} , the permutation τ covers at most two elements, namely $\tau_{[1,n-1]}$ and $\tau_{[2,n]}$. If these two elements are equal or if one of them avoids σ , then (σ, τ) contains a unique element of length $n - 1$, which contradicts the fact that (σ, τ) is disconnected. Thus, we have $\sigma < \tau_{[1,n-1]}$, $\sigma < \tau_{[2,n]}$, and $\tau_{[1,n-1]} \neq \tau_{[2,n]}$.

In \mathcal{P} , the permutation $\tau_{[2,n-1]}$ is covered by both $\tau_{[1,n-1]}$ and $\tau_{[2,n]}$. Note that $\tau_{[2,n-1]} \neq \sigma$, because $|\tau_{[2,n-1]}| = n - 2 > |\sigma|$. If $\tau_{[2,n-1]} \in (\sigma, \tau)$, then the interval (σ, τ) would be connected, because from every element there would be a path in the Hasse diagram to either $\tau_{[1,n-1]}$ or $\tau_{[2,n]}$, and thus to $\tau_{[2,n-1]}$. Therefore, $\tau_{[2,n-1]} \notin (\sigma, \tau)$, which implies that $\tau_{[2,n-1]}$ avoids σ . Together with the fact that σ is contained in both $\tau_{[1,n-1]}$ and $\tau_{[2,n]}$, it follows that σ straddles τ . \square

As we will observe in the next section, an interval is non-shellable if it contains a non-trivial disconnected subinterval. Therefore the following direct consequence of Theorem 3.2 is important to the classification of shellable intervals.

Corollary 3.3. An interval $[\sigma, \tau]$ contains a non-trivial disconnected subinterval if and only if there exists $\pi \in [\sigma, \tau]$ such that there are two adjacent occurrences of π in τ that are offset from each other by at least 3 positions. Specifically, if $\tau_i \dots \tau_j$ is the minimal consecutive subsequence of τ containing these two adjacent occurrences, then $[\pi, \tau_{[i,j]}]$ is disconnected.

For example, $[1, 2143576]$ is not itself disconnected, but the occurrences of $\pi = 21$ in 2143576 start at positions 1, 3 and 6, and we can check that $[\pi, \tau_{[3,7]}] = [21, 21354]$ is disconnected.

In the rest of the paper, we will use the term *disconnected subintervals* to mean non-trivial disconnected subintervals.

3.2. Finding disconnected intervals. Finding a disconnected subinterval using Corollary 3.3 may require checking all possible $\pi \in [\sigma, \tau]$. Our next result implies that, in some cases, it is sufficient to search in the interval $[x(\tau), \tau]$, which may contain significantly fewer elements than $[\sigma, \tau]$ (as in Example 3.5 below).

Proposition 3.4. *Suppose that $|x(\tau)| \neq 2$. If $[1, \tau]$ contains a disconnected subinterval, then so does $[x(\tau), \tau]$.*

Note that if $[\sigma, \tau]$ contains a disconnected subinterval, then so does $[1, \tau]$, and we can apply the above proposition. When $\sigma \leq x(\tau)$, this result nicely complements Theorem 2.1, which says that computing $\mu(\sigma, \tau)$ often boils down to computing $\mu(\sigma, x(\tau))$. So, when determining the Möbius function, the interesting part of the interval is often the part below $x(\tau)$, but when looking for disconnected subintervals, the interesting part is often above $x(\tau)$.

Proof. The statement is trivially true if $x(\tau) = 1$, so let us assume that $|x(\tau)| \geq 3$. Let $n = |\tau|$ and $k = |x(\tau)|$.

Suppose that $[x(\tau), \tau]$ contains no disconnected subintervals. Then, by Corollary 3.3, any two adjacent occurrences of $x(\tau)$ in τ must be offset by one or two positions. If $x(\tau)$ were monotone, then this would force τ to be monotone as well, but we know that in this case $[1, \tau]$ is a chain, which contains no disconnected subintervals.

Thus, we will assume that $x(\tau)$ is not monotone. By Lemma 2.2, no two occurrences of $x(\tau)$ in τ can be offset by one, so all adjacent occurrences of $x(\tau)$ in τ must be offset by two positions. Equivalently, $\tau_{[j,j+k-1]} = x(\tau)$ for every odd j with $1 \leq j \leq n - k + 1$, and note that n and k must have the same parity. This implies that for every i , the order relationship between τ_i and τ_{i+1} is the same as between τ_{i+2} and τ_{i+3} . By considering the complement of τ if necessary, we will assume without loss of generality that $\tau_1 < \tau_2$, and thus $\tau_i < \tau_{i+1}$ for every odd i . Since τ is not monotone, we must then have $\tau_i > \tau_{i+1}$ for every even i , and so τ is an up-down permutation, that is, $\tau_1 < \tau_2 > \tau_3 < \tau_4 > \dots$.

By Corollary 3.3, the only way for $[1, \tau]$ but not $[x(\tau), \tau]$ to contain a disconnected subinterval would be if there exists a permutation π with $x(\tau) \not\leq \pi \leq \tau$ which has two adjacent occurrences in τ offset by 3 or more. Note that we must also have $\pi \not\leq x(\tau)$, since otherwise the offsets between adjacent occurrences of π would be two. Thus, since π is incomparable with $\tau_{[j,j+k-1]}$ for every odd j , we must have that $|\pi| = k$ and any occurrence of π in τ is of the form $\tau_{[i,i+k-1]}$ for some even i . Let $a < b$ be even positions so that $\tau_{[a,a+k-1]}$ and $\tau_{[b,b+k-1]}$ are two adjacent occurrences of π offset by 3 or more, that is, π straddles $\tau_{[a,b+k-1]}$.

Suppose first that $\tau_1 < \tau_3$. Since $k \geq 3$, the fact that the patterns $\tau_{[j,j+k-1]}$ are equal for every odd j implies that $\tau_1 < \tau_3 < \tau_5 < \tau_7 < \dots$, that is, the valleys of τ are increasing. Let $\sigma = 1 \oplus \pi$. Since τ_{a-1} and τ_{b-1} are valleys, we have that $\tau_{[a-1,a+k-1]} = \tau_{[b-1,b+k-1]} = \sigma$. It follows that σ straddles $\tau_{[a-1,b+k-1]}$, since this latter permutation cannot have a third occurrence of σ without $\tau_{[a,b+k-1]}$ having a third occurrence of π . Additionally, $x(\tau) = \tau_{[a-1,a+k-2]} \leq \sigma$, since $a - 1$ is odd. Therefore, $[\sigma, \tau_{[a-1,b+k-1]}]$ is a disconnected subinterval contained in $[x(\tau), \tau]$, contradicting our original assumption.

Now consider the case $\tau_1 > \tau_3$. We now have that the valleys of τ are decreasing. If k is odd, then so is n , and we can instead consider the reversal of τ , whose first three entries form the pattern 132, and apply the above argument to it. If k is even, then $k \geq 4$, and so the order relationship between τ_2 and τ_4 determines whether the peaks of τ are increasing or decreasing. If they are increasing, then τ is the unique up-down permutation of length n with decreasing valleys and increasing peaks, and $[1, \tau]$ does not have disconnected subintervals since any adjacent occurrences of a pattern π will be offset by at most two positions. (That case will be discussed in detail below when studying the permutations with $|x(\tau)| = |\tau| - 2$.) If the peaks of τ are decreasing, then the pattern $\tau_{[i, i+k-1]}$ is the same for every even i : it starts with the largest entry k , followed by the pattern determined by the first $k-1$ entries of $x(\tau)$. It follows that adjacent occurrences of π are offset by 2 and not by 3, which is a contradiction. \square

Example 3.5. Consider the interval $[\sigma, \tau] = [1, 68372514]$. We see that $x(\tau) = 2413$, so by Proposition 3.4, we only need to look at $[2413, 68372514]$ when searching for disconnected intervals in $[\sigma, \tau]$. Using Corollary 3.3 and its notation, the only non-trivial candidates for π are 2413, 35241 and 52413. In each case, the adjacent occurrences of π in τ are offset by just two positions, so $[\sigma, \tau]$ has no disconnected subintervals. To use Corollary 3.3 without applying Proposition 3.4, we would have had eight non-trivial candidates for π to check.

The condition $|x(\tau)| \neq 2$ is used in the proof of Proposition 3.4 to argue that all adjacent occurrences of $x(\tau)$ are offset by two positions. To see an example of how the statement may fail when $|x(\tau)| = 2$, take $\tau = 1325746$. Its exterior is $x(\tau) = 12$, but the only disconnected subinterval in $[1, \tau]$ is $[21, 21453]$, which is not contained in $[x(\tau), \tau]$.

An immediate consequence of Proposition 3.4 is that $[1, \tau]$ has no (non-trivial) disconnected subintervals if $|\tau| - |x(\tau)| \leq 2$. In fact, we can precisely describe the structure of any interval $[\sigma, \tau]$ where τ has this property, as follows.

When $|\tau| - |x(\tau)| = 1$, Lemma 2.2 implies that τ is monotone, and in this case $[\sigma, \tau]$ is always a chain.

When $|\tau| - |x(\tau)| = 2$, we have that $x(\tau) = \tau_{[1, n-2]} = \tau_{[3, n]}$, where $n = |\tau|$. It follows that


$$\tau_{[i, j]} = \tau_{[i+2, j+2]} \quad \text{for every } 1 \leq i \leq j \leq n-2. \quad (3.1)$$

Thus, for every $2 \leq k \leq n-1$, there are exactly two patterns of length k contained in τ , namely $\tau_{[1, k]}$ and $\tau_{[2, k+1]}$. These patterns are different from each other, because otherwise τ would be monotone, implying that $|\tau| - |x(\tau)| = 1$. For any $\sigma \leq \tau$, the interval $[\sigma, \tau]$ has two elements of each length k with $|\sigma| < k < |\tau|$. Additionally, any two elements in $[\sigma, \tau]$ of different lengths are comparable. The Hasse diagram of such an interval is drawn in Figure 3.1. Note that the two elements of length $n-2$ are $x(\tau)$ and $i(\tau)$.

We conclude this section with a result that states that, in a certain precise sense, almost all intervals in \mathcal{P} contain disconnected subintervals. The approach below follows that in [MS15], where the analogous result is shown for the classical pattern poset.

Theorem 3.6. *Given a permutation σ , let*

$$\mathbb{P}_n([\sigma, \tau] \text{ contains a disconnected subinterval})$$

FIGURE 3.1. The interval $[\sigma, \tau]$ in the case $|\tau| - |x(\tau)| = 2$.

denote the probability that $[\sigma, \tau]$ contains a non-trivial disconnected subinterval, where τ is chosen uniformly at random from \mathcal{S}_n . Then

$$\lim_{n \rightarrow \infty} \mathbb{P}_n([\sigma, \tau] \text{ contains a disconnected subinterval}) = 1.$$

Before proving Theorem 3.6, let us give a quick preliminary result which will also be useful in Subsection 6.3.

Lemma 3.7. *Given a permutation σ , let $\mathbb{P}_n(\sigma \leq \tau)$ denote the probability that $\sigma \leq \tau$, where τ is chosen uniformly at random from \mathcal{S}_n . Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(\sigma \leq \tau) = 1.$$

Proof. Let $k = |\sigma|$. Let us break the first $k\lfloor n/k \rfloor$ entries of τ into $\lfloor n/k \rfloor$ disjoint blocks of length k , starting at the beginning. The probability that none of these blocks is an occurrence of σ is

$$\left(1 - \frac{1}{k!}\right)^{\lfloor n/k \rfloor}.$$

Since k is fixed, this probability clearly goes to 0 as $n \rightarrow \infty$, implying the desired result. \square

Proof of Theorem 3.6. Letting $k = |\sigma|$, we can assume that $k \geq 3$, since establishing the result for such permutations will also clearly prove it for all σ . Assume first that $\sigma_1 > \sigma_k$. In this case, we see that σ straddles $\sigma \oplus \sigma$ and so, by Theorem 3.2, $[\sigma, \sigma \oplus \sigma]$ is disconnected. Applying Lemma 3.7 tells us that τ contains $\sigma \oplus \sigma$ almost surely as $n \rightarrow \infty$. We conclude that as $n \rightarrow \infty$, the interval $[\sigma, \tau]$ contains the disconnected interval $[\sigma, \sigma \oplus \sigma]$ with probability approaching 1, as required.

If $\sigma_1 < \sigma_k$, then σ straddles $\sigma \ominus \sigma$ and we proceed similarly. \square

In contrast with Theorem 3.6, we note that there are infinite classes of intervals that do not contain disconnected subintervals. For a simple example, any time σ satisfies $\sigma_1 > \sigma_{|\sigma|}$, we can take monotone increasing sequences α and β and let $\tau = \alpha \oplus \sigma \oplus \beta$. Then any element of $[\sigma, \tau]$ will occur only once in τ , so there are no disconnected subintervals. For examples of intervals $[\sigma, \tau]$ that are not just a

product of two chains but contain no disconnected subintervals, let $\sigma = 21$ and $\tau = \alpha \oplus 21 \oplus 21 \oplus \cdots \oplus 21 \oplus \beta$, where the number of copies of 21 is at least 2 and where α and β are monotone increasing as before.

4. SHELLABILITY

Our goal for this section is to prove that all intervals in \mathcal{P} are shellable, except for those which are not shellable for a straightforward reason. We begin by giving the necessary background on shellability. We refer the interested reader to [Wac07] for further details and a wealth of other information about poset topology.

4.1. Background on shellability and CL-shellability. When considering a combinatorially defined simplicial complex, such as the order complex $\Delta(\sigma, \tau)$ of an interval $[\sigma, \tau]$ of \mathcal{P} , it is common to ask if the simplicial complex is shellable. Since our intervals $[\sigma, \tau]$ are graded, we can restrict our discussion to order complexes that are pure, i.e., all the facets have the same dimension. A *shelling* of a pure d -dimensional simplicial complex Δ is a linear ordering F_1, F_2, \dots, F_s of its facets such that the intersection

$$\left(\bigcup_{j=1}^{k-1} F_j \right) \cap F_k \tag{4.1}$$

is pure and $(d-1)$ -dimensional for all $k = 2, 3, \dots, s$. A simplicial complex is *shellable* if it has a shelling. Figure 2.1 is an example of a non-shellable complex since in any ordering of the facets, the point 213 will be a $(d-2)$ -dimensional connected component of an intersection of the type in (4.1), implying that no shelling exists. As is customary, we will say that $[\sigma, \tau]$ is *shellable* if $\Delta(\sigma, \tau)$ is.

Much of the interest in shellability arises from the fact that if an interval $[\sigma, \tau]$ is shellable, then this tells us that the homotopy type of $\Delta(\sigma, \tau)$ is a wedge of spheres of the top dimension. The number of spheres is $|\mu(\sigma, \tau)|$ which, in the case of intervals in \mathcal{P} , can be determined from Theorem 2.1. In the shellable case, we deduce that either the number of spheres is one or $\Delta(\sigma, \tau)$ is contractible.

These topological considerations are also why we take the permutation 1 as the bottom element of \mathcal{P} , rather than the empty permutation \emptyset : since \emptyset is only covered by the permutation 1, any complex of the form $\Delta(\emptyset, \pi)$ will be contractible.

We will prove shellability using Björner and Wachs's theory of CL-shellability, which is a generalization¹ of Björner's theory of EL-shellability [Bjö80]. The original definition of CL-labeling was given in [BW82] under the name "L-labeling"; the now customary definition below appears in [BW83]. We will partially follow [Wac07] in the exposition here.

A poset Q is said to be *bounded* if it has a unique minimal and a unique maximal element, which we denote $\hat{0}$ and $\hat{1}$ respectively. When x is covered by y , we will write the corresponding edge as $x \leftarrow y$; the head of the arrow reflects the $<$ symbol, and this notation will seem natural later when we primarily follow chains from top to bottom. Rather than just labeling edges $x \leftarrow y$ as is the case for EL-labelings, we consider the set $\mathcal{ME}(Q)$ of pairs $(C, x \leftarrow y)$ consisting of a maximal chain C and an edge along that chain. A *chain-edge labeling* of Q is a map $\lambda : \mathcal{ME}(Q) \rightarrow \Lambda$, where Λ is some poset, satisfying the following condition: if two maximal chains coincide

¹Actually, it remains open whether there exist posets that are CL-shellable but not EL-shellable. The point of CL-shellability is that it is more flexible than EL-shellability, making CL-shellings sometimes easier to find.




FIGURE 4.1. A chain-edge labeling that is a CL-labeling. Notice that the edge $d \leftarrow f$ receives two possible labels: the label 2 as an element of the chain $a \leftarrow b \leftarrow d \leftarrow f$, and the label 1 as an element of the chain $a \leftarrow c \leftarrow d \leftarrow f$. The bottom-rooted interval $[b, f]_{a \leftarrow b}$ has a unique increasing chain, namely $b \leftarrow d \leftarrow f$ which has label sequence (2, 2).

along their bottom k edges, then their labels also coincide along these edges. The key point is that the label on an edge depends on the maximal chain along which we arrive at that edge. See Figure 4.1 for an example.

It follows that to restrict a chain-edge labeling of Q to an interval $[u, v]$ in Q , we need to record how we arrived at the bottom element u of the interval. Therefore, we introduce the idea of a *bottom-rooted interval* $[u, v]_r$, which is the interval $[u, v]$ together with a maximal chain r of $[\hat{0}, u]$. This rooting allows the restriction of λ to an interval $[u, v]$ to be consistent with λ defined over all of Q . More precisely, for an edge $x \leftarrow y$ in $[u, v]_r$ the label received by $x \leftarrow y$ when considered as being along the maximal chain C of $[u, v]_r$ is equal to the label received by $x \leftarrow y$ when considered as being along the maximal chain $r \cup C \cup s$ of Q , where s is any maximal chain of $[v, \hat{1}]$.

When we say that a maximal chain M of $[u, v]_r$ is *weakly increasing* or *lexicographically precedes* another, we are referring to the sequence of labels along M as we read from u up to v .

Definition 4.1. Let Q be a bounded poset. A chain-edge labeling λ is a *CL-labeling* (chain-lexicographic labeling) if in each bottom-rooted interval $[u, v]_r$ of Q , there is a unique weakly increasing maximal chain, and this chain lexicographically precedes all other maximal chains in $[u, v]_r$. A poset which has a CL-labeling is said to be *CL-shellable*.

An example of a CL-labeling is shown in Figure 4.1. As the terminology suggests, a bounded poset that is CL-shellable is shellable [BW82, Theorem 3.3] and, along with EL-shellability, exhibiting CL-shellability is one of the most common ways to prove shellability of order complexes.

4.2. Shellability in the consecutive pattern poset. It is easy to see that disconnected intervals of rank at least 3 are not shellable. Moreover, a result of Björner [Bjö80] states that if a poset is shellable, then all of its subintervals are shellable. In particular, if an interval $[\sigma, \tau]$ contains a disconnected subinterval of rank at least

3, then it is certainly not shellable. As an example, the disconnected subinterval [213, 213546] in Figure 1.1 results in the non-shellability of the corresponding order complex in Figure 2.1.

It thus follows immediately from Theorem 3.6 that almost all intervals in \mathcal{P} are not shellable, in a particular precise sense.

Corollary 4.2. *Given a permutation σ , let τ be chosen uniformly at random from S_n . Then the probability that $[\sigma, \tau]$ is shellable tends to 0 as n tends to infinity.*

Our main result of this section is that the converse of Björner’s result is true in \mathcal{P} : if an interval $[\sigma, \tau]$ does not contain a disconnected subinterval, then it is shellable. So, roughly speaking, all intervals in the consecutive pattern poset that have any hope of being shellable are in fact shellable. Recall that we have already classified those intervals that contain disconnected subintervals in Corollary 3.3 in terms of the entries of σ and τ , so the theorem we are about to state classifies those σ and τ that yield a shellable interval $[\sigma, \tau]$.

Theorem 4.3. *The interval $[\sigma, \tau]$ in \mathcal{P} is shellable if and only if it contains no non-trivial disconnected subintervals.*

It remains to prove the “if” direction, which we will do using CL-shellability. Because we will read the labels on our maximal chains from top to bottom, we will actually show that the dual of $[\sigma, \tau]$ is CL-shellable and hence shellable; this implies the shellability of $[\sigma, \tau]$ since the order complex of $[\sigma, \tau]$ is clearly isomorphic to that of its dual. The analogue of Theorem 4.3 in the classical pattern poset is actually false ([123, 3416725] is given in [MS15] as a counterexample), but we should note that some of the ideas of our proof are similar to those in [MS15] for proving the dual CL-shellability in the classical case of intervals of layered permutations without non-trivial disconnected subintervals.

Before giving the proof, we will define the chain-edge labeling that we will use. We will first set provisional labels using a first pass, and then make some modifications in a second pass. If π' has length ℓ and covers π along some maximal chain C of $[\sigma, \tau]$ and $\pi = \pi'_{[2,\ell]}$, then we label the edge $\pi' \rightarrow \pi$ by 0. Otherwise, we assign the label 1. Referring to Lemma 2.2, note that we assign the label 0 whenever π' is monotone. However, the resulting labeling, which is actually an edge-labeling (i.e., each edge label is independent of the chain C on which the edge appears), is not a dual CL-labeling. For example, the disconnected rank-2 intervals [21, 2143] and [213, 21435] each have two increasing chains from top to bottom, one with labels (0,0) and the other with labels (1,1). Here and elsewhere, we say “increasing” to mean weakly increasing. We next describe the modification we make to our labeling to address this shortcoming.

Pick $\varepsilon > 0$ to be small; $\varepsilon \leq 1/(|\tau| - |\sigma|)$ will suffice. For every maximal chain C , work along C from top to bottom, pausing at any triple $\pi'' \rightarrow \pi' \rightarrow \pi$ such that π straddles π'' , as in the examples of the previous paragraph. We see that this situation of straddling for rank-2 intervals is exactly the condition that results in chains with labels (0,0) and (1,1). Do nothing to the two labels that are already 0. For the other chain, letting $\lambda_C(\pi', \pi)$ denote the label of the edge $\pi' \rightarrow \pi$ along C , decrease the label on $\pi' \rightarrow \pi$ so that $\lambda_C(\pi', \pi) = \lambda_C(\pi'', \pi') - \varepsilon$. After completion of these modifications, every label will either be 0 or $1 - k\varepsilon$ for some $k \geq 0$. See Figure 4.2 for an example. The edge 213 → 21 there shows that this labeling is a chain-edge labeling rather than just an edge labeling.




FIGURE 4.2. The interval $[21, 214356]$ with its increasing chain shown in bold. The top-rooted intervals $[213, 21435]_{214356 \rightarrow 21435}$ and $[21, 2143]_{214356 \rightarrow 21435 \rightarrow 2143}$ have had their labels modified in the particular way we describe. The edge labeled 1^* takes the label $1 - 2\varepsilon$ along the chain that arrives via 2143 and otherwise takes the label 1 . All other labels are independent of the chain on which they appear.

Proof of Theorem 4.3. As already mentioned in the first paragraph of this subsection, the interval $[\sigma, \tau]$ is not shellable if it contains a non-trivial disconnected subinterval. For the converse, we will prove dual CL-shellability. So suppose $[\sigma, \tau]$ contains no non-trivial disconnected subintervals, and let $[\alpha, \beta]_r$ be a *top-rooted interval* in $[\sigma, \tau]$, the obvious analogue of bottom-rooted interval. Using the labeling λ_C defined above and reading labels along chains from top to bottom, we wish to show that there is a unique increasing maximal chain in $[\alpha, \beta]_r$ and that this increasing chain has the lexicographically first labels of all maximal chains in $[\alpha, \beta]_r$.

Every maximal chain D from β to α can be said to “end” at a particular occurrence of α in β as follows. Each edge of D corresponds to deleting an entry from the left or from the right of β , with the convention that the outgoing edge of a monotone permutation corresponds to deleting from the left (this is consistent with the label 0 assigned in our definition of λ_C); the entries not deleted from β after traversing D give an occurrence of α .

Define a particular maximal chain D of $[\alpha, \beta]_r$ in the following way: starting at β , delete the leftmost entry if the result will still be in $[\alpha, \beta]_r$, and otherwise delete the rightmost entry. In other words, D will delete from the left of β until it reaches the rightmost occurrence of α in β , at which point it will switch to deleting from the right. None of the resulting labels will be of the form $1 - k\varepsilon$ for $k > 0$ since D will never delete from the right when there is also the option of deleting from the left; see also the technical note in the next paragraph. Thus the sequence of labels will take the form $0, 0, \dots, 0, 1, 1, \dots, 1$ from top to bottom, possibly with no 0’s or no 1’s. By definition of D , these labels are the lexicographically least of any maximal chain.

It is necessary to make one technical note at this point. One might wonder about whether the deletions along r could cause any of the labels along D to be of the form $1 - k\varepsilon$ with $k > 0$. Let us show that this never happens. We see that the only possibility is that the first label within $[\alpha, \beta]_r$ along D is of this form. By definition of D , this implies that D always deletes from the right, and so α only occurs in β as a prefix. Thus there is only one maximal chain D in $[\alpha, \beta]_r$ and it has label sequence $(1 - k\varepsilon, 1, 1, \dots, 1)$ from top to bottom, which is increasing as required. While this shows that labels of the form $1 - k\varepsilon$ with $k > 0$ on D do not cause any difficulty, it does not fulfill our promise to show that such labels never occur. So let β^+ be the element covering β on r , and let β^- be the element β covers on D . If $\beta \rightarrow \beta^-$ gets the label $1 - k\varepsilon$ with $k > 0$, this means by definition of our label modifications that $\beta^- = \beta_{[1, \ell-2]}^+ = \beta_{[3, \ell]}^+$ where $\ell = |\beta^+|$. Since α only occurs in β and hence β^- as a prefix, the element covered by β^- along D is then $\beta_{[1, \ell-3]}^+ = \beta_{[3, \ell-1]}^+$. But $\beta_{[3, \ell-1]}^+ = \beta_{[3, \ell-1]}$, contradicting the fact that α only occurs in β as a prefix. We conclude that all of the labels along D are indeed 0 or 1.

Now suppose $[\alpha, \beta]_r$ contains another increasing maximal chain D' . Our approach will be to consider the possible locations of the occurrence of α at which D' ends, and to show some contradiction in each case. If D' ends at the rightmost occurrence $\beta_i \dots \beta_j$ of α in β , then D' must equal D to be increasing, since the path along D' must make all its left deletions before making its right deletions. Therefore, assume D' ends at an occurrence of α in β that is not rightmost, namely $\beta_{i'} \dots \beta_{j'}$. Since D' is increasing, its last deletion is on the right and thus receives a label of the form $1 - k\varepsilon$ for $k \geq 0$. This already yields a contradiction in the case when $\sigma = 1$ since any permutation covering the permutation 1 is monotone. Therefore, assume that $|\sigma| > 1$ for the rest of this proof. Among those increasing maximal chains different from D , choose D' so that i' is as large as possible. There are three cases to consider.

First suppose $i' = i - 1$ which, by Lemma 2.2, implies that $\beta_{[i', j'+1]}$ is monotone. Since the last deletion along D' is on the right, the element covering α on D' is this monotone permutation $\beta_{[i', j'+1]}$. But our convention for monotone permutations then implies that the last edge in D' is labeled 0, which is a contradiction.

Next suppose that $i' = i - 2$. Thus $\beta_{[i-2, j-2]} = \beta_{[i, j]} = \alpha$. If it is also the case that $\beta_{[i-1, j-1]} = \alpha$, then applying Lemma 2.2 three times tells us that $\beta_{[i-2, j]}$ is monotone. Since the last deletion of D' is on the right, the element covering α on D' will be $\beta_{[i-2, j-1]}$, obtaining a contradiction as in the previous paragraph. Thus $\beta_{[i-1, j-1]} \neq \alpha$, and so α straddles $\beta_{[i-2, j]}$. Since D' ends at $\beta_{[i-2, j-2]}$, the last two edges on D' must receive the label $1 - k\varepsilon$ followed by the label $1 - (k+1)\varepsilon$ for some $k \geq 0$. This contradicts D' having increasing labels.

For the third and final case, let $i' < i - 2$. Suppose first that, except for $\beta_i \dots \beta_j$, the rightmost occurrence of α in β is $\beta_{i'} \dots \beta_{j'}$. Then consider the subinterval $[\alpha, \beta_{[i', j]}]$ of $[\alpha, \beta]_r$. Since α straddles $\beta_{[i', j]}$, Theorem 3.2 implies that $[\alpha, \beta]_r$ has a non-trivial disconnected subinterval, contradicting our hypothesis. Thus there must be some other occurrence $\beta_{i''} \dots \beta_{j''}$ of α with $i' < i'' < i$ and, among all such occurrences, let us choose i'' to be as close to i' as possible. If $i'' = i' + 1$, then $\beta_{[i', j'+1]}$ is monotone, yielding a contradiction as before. If $i'' = i' + 2$, then we can apply the argument of the previous paragraph with i'' and j'' in place of i and j to yield a contradiction. Finally, suppose $i'' > i' + 2$, and consider the subinterval $[\alpha, \beta_{[i', j''}]]$ of $[\alpha, \beta]_r$. By our choice of i'' , we see that α straddles $\beta_{[i', j'']}$, yielding

by Theorem 3.2 a non-trivial disconnected subinterval of $[\alpha, \beta]_r$, contradicting our hypothesis. \square

We conclude this section with a comparison of Theorem 4.3 to related results in the literature.

Remark 4.4. It follows from Theorem 4.3 that the homotopy type of an interval $[\sigma, \tau]$ without non-trivial disconnected subintervals is a wedge of spheres. As we mentioned in Subsection 4.1, the number of spheres is $|\mu(\sigma, \tau)|$ so, by Theorem 2.1, we conclude that $\Delta(\sigma, \tau)$ is homotopic to a single sphere or is contractible. This same conclusion is given as [SW12, Theorem 2.8] but they do not show shellability and instead use discrete Morse theory. In fact, their conclusion applies to all intervals $[\sigma, \tau]$ in \mathcal{P} , even those containing disconnected subintervals. An example of this more general setting is the 2-dimensional simplicial complex in Figure 2.1, which contracts to a sphere of dimension 1.

It also follows from Theorem 4.3 that the single spheres that arise must be of the same dimension as $\Delta(\sigma, \tau)$, but this is no longer true in the non-shellable cases, as shown by the aforementioned example in Figure 2.1. In summary, Theorem 4.3 has a stronger hypothesis but also a stronger conclusion than [SW12, Theorem 2.8].

Remark 4.5. Billera and Myers [BM99] have shown a general result of a similar flavor to Theorem 4.3: a bounded graded poset is shellable if it is $(2+2)$ -free, meaning that it contains no induced subposet that is isomorphic to a disjoint sum of two chains of length 2. In [Wac99], Wachs shows that such $(2+2)$ -free posets are CL-shellable. We note, however, that these results are not enough to imply Theorem 4.3, since there exist intervals in \mathcal{P} that contain no non-trivial disconnected subintervals but fail to be $(2+2)$ -free. One example is in Figure 4.2, where the two chains $2143 < 21435$ and $2134 < 13245$ are of the form $2 + 2$.

5. ALL INTERVALS ARE RANK-UNIMODAL AND STRONGLY SPERNER

Since intervals in \mathcal{P} are graded, it is natural to ask about the structure of these intervals with regard to the number of elements at each rank. It is clear from Figure 1.1 that the intervals are not rank-symmetric in general. It is also the case that the sequence of rank sizes is not log-concave in general. For example, the interval $[1, 1265473]$ has one element of rank 0 and five elements of rank 2 but only two elements of rank 1. However, there are two interesting properties that intervals of \mathcal{P} have: they are rank-unimodal and strongly Sperner. Our goal in this section is to prove this assertion. We note that neither result is known for the classical pattern poset [MS15]. Even though the definitions of rank-unimodal and strongly Sperner are quite different, the two results appear together because they both rely heavily on the same injection between rank levels (Lemma 5.1), and because Theorem 5.5 below of Griggs connects the two properties.

5.1. Rank-unimodality. For a finite graded poset Q , the set of elements of rank i will be called a *rank level* of Q , and the cardinality a_i of this rank level will be referred to as the *size* of rank i . Recall that a finite graded poset of rank N is called *rank-unimodal* if the sequence a_0, a_1, \dots, a_N of rank sizes is unimodal.

Before proving rank-unimodality for intervals in \mathcal{P} , it will be helpful and informative to examine an explicit 3-step method for constructing such intervals, up to isomorphism. Fix $\sigma \leq \tau$, and let $n = |\tau|$. Consider the poset P_1 of intervals

$[i, j] = \{i, i+1, \dots, j\}$ with $1 \leq i \leq j \leq n$, ordered by set inclusion. Each interval $[i, j]$ is associated with a subsequence $\tau_i \dots \tau_j$ of τ . The poset P_1 is a graded join-semilattice of rank $n-1$, having $n-r$ elements or rank r for all r , namely $[i, i+r]$ for $1 \leq i \leq n-r$. See Figure 5.1 for an example.




FIGURE 5.1. For $[\sigma, \tau] = [12, 213546]$, the full poset above shows P_1 , with P_2 given by the elements outlined with bold lines. Figure 5.2 shows P_3 . In this example, the breaking rank is $b = 1$ because $[4, 6] \sim [1, 3]$.

For the second step, let P_2 be the poset obtained from P_1 by deleting all the elements $[i, j]$ such that $\tau_{[i,j]}$ avoids σ . For the final step, let P_3 be the quotient poset² P_2 / \sim , where \sim is the equivalence relation defined by $[i, j] \sim [i', j']$ whenever $\tau_{[i,j]} = \tau_{[i',j']}$. For equivalence classes in the quotient poset, denoted $\overline{[i,j]}$, recall that one defines $\overline{[i,j]} \leq \overline{[i',j']}$ if there exist $[i_1, j_1] \in \overline{[i,j]}$ and $[i_2, j_2] \in \overline{[i',j']}$ such that $[i_1, j_1] \subseteq [i_2, j_2]$. It is clear that the map $\overline{[i,j]} \mapsto \tau_{[i,j]}$ is a poset isomorphism between P_3 and the interval $[\sigma, \tau]$ in \mathcal{P} ; see Figure 5.2 for an example. In what follows, we will work with P_3 instead of $[\sigma, \tau]$ when helpful.

From this construction of $[\sigma, \tau]$, we can make our first observation of this section about the structure of $[\sigma, \tau]$: the upper ranks of $[\sigma, \tau]$ are “grid-like.” More precisely, let $N = |\tau| - |\sigma|$ denote the rank of $[\sigma, \tau]$ (or of P_3). By comparing the structure of P_1 and P_3 , we know there are at most $N + 1 - r$ elements of rank r in $[\sigma, \tau]$. Let b denote the largest value of r such that the number of elements of rank r is strictly less than $N + 1 - r$. We will call b the *breaking rank* of $[\sigma, \tau]$. In the context of P_3 , b is the rank of $\overline{[i, i+j]}$ where j is the largest value such that either an element of the form $[i, i+j]$ was deleted when going from P_1 to P_2 , or there is an equivalence class of the form $\overline{[i, i+j]}$ in P_3 containing more than one element of P_2 . For example, for $[12, 213546]$ exhibited in Figures 5.1 and 5.2, we have $b = 1$ since $[4, 6] \sim [1, 3]$. Notice that these differences between P_1 and P_3 that happen at the breaking rank

² There is more than one notion of quotient poset in the literature. Our construction follows Hallam and Sagan [HS15]: given an equivalence relation \sim on a poset P , the quotient P / \sim is the set of equivalence classes, with $X \leq Y$ in P / \sim if and only if $x \leq y$ in P for some $x \in X$ and some $y \in Y$. Hallam and Sagan show that P / \sim is a poset if whenever $X \leq Y$ in P / \sim , we have that for all $y \in Y$ there exists $x \in X$ such that $x \leq y$ in P . We see that P_2 / \sim satisfies this condition and so P_3 is indeed a poset.




FIGURE 5.2. The interval $[12, 213546]$ in \mathcal{P} has already appeared in Figure 1.1, but here it is adorned with the corresponding elements of P_3 . The bold edges show two examples of the injection of Lemma 5.1. For the injection from $m_1 = 2$ elements of E_1 into E_2 , we chose $k = 2$ in the notation of the proof of Lemma 5.1. For the injection from $m_2 = 2$ elements of E_2 into E_3 , our only valid option is $k = 3$.

trickle down P_3 , meaning that if r is a rank less than the breaking rank then there will again be strictly less than the maximum $N + 1 - r$ elements of rank r .

By the “grid-like” structure of the upper ranks of $[\sigma, \tau]$ we mean that, by definition of the breaking rank b , the rank levels of rank r with $b < r \leq N$ of $[\sigma, \tau]$ are essentially the same as the top $N - b$ rank levels of P_1 . For example, in Figure 5.1, the grid-like form of the top three rank levels gives rise to the grid-like form of the top three rank levels in Figure 5.2. This observation about the upper ranks directly relates to the rank-unimodality of $[\sigma, \tau]$ since it implies

$$a_{b+1} > a_{b+2} > \cdots > a_N, \quad (5.1)$$

which will be the basis of the “downward portion” of the unimodality.

The following lemma will be used in both the proof of rank-unimodality and of the strongly Sperner property, and allows us to pair up elements of adjacent ranks.

Lemma 5.1. *In the interval $[\sigma, \tau]$ in \mathcal{P} , let E_r denote the set of elements of rank r for $0 \leq r \leq N$, where $N = |\tau| - |\sigma|$. For $0 \leq r \leq N - 1$, pick any $m_r \leq \min\{|E_r|, N - r\}$. Then for any selection of m_r elements of E_r , there exists an injection f_r from these elements into E_{r+1} with the property that $\pi < f_r(\pi)$ for each such element π .*

Proof. We will work in the setting of P_3 as defined above, and the result will follow by the isomorphism with $[\sigma, \tau]$. Fix m_r elements of rank r . These elements of rank r in P_3 take the form $[i, i+s]$, where $s = |\sigma| - 1 + r$. For each such element, choose the representative $[i, i+s]$ with smallest i . Let $I_s \subset \{1, 2, \dots, n-s\}$ be the set of left endpoints of these m_r representatives, where $n = |\tau|$ as before. Since $m_r \leq N - r = n - s - 1$, there is some $k \in \{1, 2, \dots, n-s\} \setminus I_s$. Fix such a k , and

define, for each $i \in I_s$,

$$f(i) = \begin{cases} i & \text{if } i < k, \\ i-1 & \text{if } i > k. \end{cases}$$

We claim that the map $\overline{[i, i+s]} \mapsto \overline{[f(i), f(i)+s+1]}$, defined for $i \in I_s$, is an injection from elements of rank r to elements of rank $r+1$ in P_3 . See Figure 5.2 for examples.

To see that this map is well defined, note first that $1 \leq f(i) \leq n-s-1$ for all $i \in I_s$, and so $\tau_{[f(i), f(i)+s+1]}$ is defined. Additionally, both $\tau_{[i, i+s+1]}$ and $\tau_{[i-1, i+s]}$ (when defined) contain $\tau_{[i, i+s]}$, which translates to the desired property that $\pi < f_r(\pi)$ for all π . Since $\tau_{[i, i+s]}$ in turn contains σ , we get that $\overline{[f(i), f(i)+s+1]}$ is indeed an element of P_3 .

Now we show that the map is injective. Clearly, if $[i, i+s] \sim [i', i'+s]$, then $[i-1, i+s] \sim [i'-1, i'+s]$ and $[i, i+s+1] \sim [i', i'+s+1]$. Thus, the only possible way that injectivity could fail would be if there exist $i, i' \in I_s$ with $i < k < i'$ such that $[i, i+s+1] \sim [i'-1, i'+s]$. But then $[i+1, i+s+1] \sim [i', i'+s]$, contradicting the choice of $i' \in I_s$ as the smallest left endpoint of a representative of the class. \square

We can now deduce a combinatorial proof of the rank-unimodality of all intervals in \mathcal{P} .

Corollary 5.2. *Every interval $[\sigma, \tau]$ in \mathcal{P} is rank-unimodal.*

Proof. Since we have already shown (5.1), it remains to show that $a_0 \leq a_1 \leq \dots \leq a_b \leq a_{b+1}$. To see this, fix r with $0 \leq r \leq b$ and consider the elements of rank r in $[\sigma, \tau]$. By definition of the breaking rank b , there are strictly less than $N+1-r$ elements of rank r . The result now follows directly from the injection of Lemma 5.1. \square

Remark 5.3. In the case of the classical pattern poset, it is an open question whether every interval is rank-unimodal [MS15]. It is known to be true for intervals $[\sigma, \tau]$ with $|\tau| \leq 8$.

5.2. Strongly Sperner property. This subsection is devoted to a proof that all intervals in \mathcal{P} are strongly Sperner. Let us first give the necessary background. Recall that a graded poset Q is said to be *Sperner* if the largest rank size equals the size of the largest antichain. Since every rank level of Q consists of an antichain, the Sperner property is equivalent to the condition that some rank level of Q is an antichain of maximum size. For any positive integer k , a *k-family* of Q is a union of k antichains. We are now ready for the definition of one of the well-known generalizations of the Sperner property.

Definition 5.4. Let Q be a poset of rank N . For an integer k with $1 \leq k \leq N+1$, we say that Q is *k-Sperner* if the sum of the sizes of the k largest ranks equals the size of the largest k -family. In addition, Q is said to be *strongly Sperner* if it is k -Sperner for all k .

Again, being k -Sperner is equivalent to having a set of k ranks that is a k -family of maximum size. Our technique for showing that every interval of \mathcal{P} is strongly Sperner will make heavy use of the injections f_r from Lemma 5.1 and of Theorem 5.5 below, due to Griggs [Gri80], with shorter proofs appearing in [GSS80, Lec97].

Letting Q be a poset of rank N , for all $1 \leq i \leq N + 1$, let us denote the i th largest rank level by L_i and its size by ℓ_i , where we break ties in size arbitrarily. For example, for the interval $[12, 213546]$ of Figure 5.2, we have two choices for L_1 , both of which give $\ell_1 = 3$, while $\ell_4 = \ell_5 = 1$. We will say that Q is *i -rank intersecting* if there exist ℓ_i disjoint chains which intersect each of the rank levels L_1, L_2, \dots, L_i . In particular, note that these disjoint chains must cover L_i . In addition, we say that Q is *strongly rank intersecting* if it is i -rank intersecting for all i with $1 \leq i \leq N + 1$.

Theorem 5.5 ([Gri80]). *If a poset Q is rank-unimodal, then Q is strongly Sperner if and only if it is strongly rank intersecting.*

We now have the tools necessary to prove the promised result.

Theorem 5.6. *Every interval $[\sigma, \tau]$ in \mathcal{P} is strongly Sperner.*

Proof. By Theorem 5.5, it suffices to show that $Q = [\sigma, \tau]$ is strongly rank intersecting, and throughout this proof we will follow the notation of the paragraph immediately preceding Theorem 5.5. Thus, fixing i with $1 \leq i \leq N + 1$, consider the i largest rank levels L_1, L_2, \dots, L_i . Since $[\sigma, \tau]$ is rank-unimodal and we can break ties among rank sizes arbitrarily, we can assume that L_1, L_2, \dots, L_i form a set of consecutive rank levels in some order, with L_i occurring at one of the two extremes of the consecutive sequence. Let us suppose that these rank levels occupy ranks r_1 through r_2 in $[\sigma, \tau]$.

As we will see, the injections f_r of Lemma 5.1 will be exactly what we need to complete the proof. Since we wish to construct ℓ_i disjoint chains, we will be picking ℓ_i elements of rank r for all $r_1 \leq r < r_2$. In order to apply the lemma, we need to show that $\ell_i \leq N - r$ for all such r , meaning that we need $\ell_i \leq N - r_2 + 1$. By definition of L_i and hence ℓ_i , we have $\ell_i = \min\{a_{r_1}, a_{r_2}\}$, where a_j again denotes the number of elements of rank j . We already showed when considering the grid-like structure of the upper levels of $[\sigma, \tau]$ that $a_r \leq N + 1 - r$. Thus $\ell_i \leq a_{r_2} \leq N - r_2 + 1$, as required.

Since $\ell_i = \min\{a_{r_1}, a_{r_2}\}$, we can start by picking ℓ_i elements of rank r_1 . Applying f_{r_1} to these elements, we construct ℓ_i disjoint chains of length 1 from rank r_1 to rank $r_1 + 1$. Continuing in this fashion, applying $f_{r_1+1}, f_{r_1+2}, \dots, f_{r_2-1}$, we construct ℓ_i disjoint chains that intersect each of the rank levels r_1 through r_2 , as required. See the bold edges in Figure 5.2 for an example with $i = 3$. \square

Remark 5.7. Generalizing the Sperner property in a different way, a poset is said to be *strictly Sperner* if every antichain of maximum size is a rank level. To see that this property does not hold for intervals in \mathcal{P} , consider the antichain $\{123, 1432\}$ in the interval $[12, 12543]$.

We conclude this section with an open problem. When σ occurs just once in τ , we know from Proposition 2.4 that $[\sigma, \tau]$ is a lattice.

Problem 5.8. *Characterize those intervals $[\sigma, \tau]$ in \mathcal{P} that are lattices. As usual, we would prefer our characterization to be in terms of simple conditions on σ and τ .*

For example, for $|\tau| = 4$, $[1, \tau]$ is a lattice if and only if either $\tau_{[1,3]}$ or $\tau_{[2,4]}$ (or both) is monotone, since otherwise both will be upper bounds for 12 and 21.

6. THE EXTERIOR OF A PERMUTATION

6.1. The length of the exterior. As we have seen in the previous sections, particularly in Theorem 2.1 and Proposition 3.4, the exterior of a permutation plays an important role in the study of \mathcal{P} . From an enumerative perspective, a natural problem is to study the distribution of the statistic *length of the exterior* on permutations. Table 6.1 shows the number of permutations according to this statistic. The first column (after the one indexing n) counts permutations whose exterior has length 1. Permutations with this property are usually called *non-overlapping permutations* in the literature. It was shown by Bóna [Bón11] that the number of non-overlapping permutations of length n is approximately $0.364n!$, but no exact formula is known. Next we show a simple divisibility property of these numbers.

$n \setminus k$	1	2	3	4	5	6	7	8	9
2	2								
3		4	2						
4		12	10	2					
5		48	58	12	2				
6		280	306	118	14	2			
7		1864	2186	822	150	16	2		
8		14840	17034	6580	1660	186	18	2	
9		132276	154162	58854	15118	2222	226	20	2
10		1323504	1532574	588898	150388	30238	2904	270	22

TABLE 6.1. The number of permutations $\tau \in \mathcal{S}_n$ with $|x(\tau)| = k$.

Lemma 6.1. *For every $n \geq 3$, the number of non-overlapping permutations in \mathcal{S}_n is divisible by 4.*

Proof. If $\tau \in \mathcal{S}_n$ is non-overlapping, then so are its reversal, its complement, and its reverse-complement, and these four permutations are all different. Indeed, if two of these were the same, then τ would be equal to its reverse-complement, which would imply that τ starts with a descent if and only if it ends with a descent, and so τ would be overlapping. Thus, we have partitioned the set of non-overlapping permutations into disjoint sets of size 4. \square

While the above property of the first column of Table 6.1 is straightforward to prove, there seem to be other less obvious congruence relations. For example, the data in the second column suggests that the number of permutations in \mathcal{S}_n with exterior of length 2 might always be congruent to 2 modulo 4.

The entries in the rightmost nonzero diagonal of Table 6.1 count permutations τ with $|x(\tau)| = |\tau| - 1$. It is clear that the number of such permutations in \mathcal{S}_n is 2 for every $n \geq 2$, since by Lemma 2.2, only the monotone permutations $12\dots n$ and $n\dots 21$ satisfy this condition. The entries in the diagonal immediately below these 2s count permutations τ with $|x(\tau)| = |\tau| - 2$. In Subsection 3.2 we described the structure of the intervals under such permutations. Next we enumerate these permutations.

Lemma 6.2. *For every $n \geq 4$,*

$$|\{\tau \in \mathcal{S}_n : |x(\tau)| = n - 2\}| = 2n + 2.$$

Proof. Since the length of the exterior of any permutation is the same for its complement, we can count permutations with $\tau_1 < \tau_2$ and multiply the final number by 2. We have shown that every $\tau \in \mathcal{S}_n$ with $|x(\tau)| = n - 2$ satisfies Equation (3.1). It follows that $\tau_i < \tau_{i+1}$ for every odd i and, since τ is not monotone (otherwise $|x(\tau)| = n - 1$), that $\tau_i > \tau_{i+1}$ for every even i , so τ is alternating.

By Equation (3.1), the relative order of τ_1 and τ_3 determines whether the sequence of valleys of τ is increasing or decreasing, and similarly the relative order of τ_2 and τ_4 determines it for the sequence of peaks. When one of these sequences is increasing and the other decreasing, which happens when $\tau_{[1,4]}$ equals 1423 or 2314, then τ is completely determined.

If both sequences are increasing, that is, $\tau_{[1,4]} = 1324$, then τ is determined once we choose the value of τ_2 relative to the valleys of τ , that is, for which index $2 \leq i \leq \lfloor \frac{n+1}{2} \rfloor$ we have $\tau_{2i-1} < \tau_2 < \tau_{2i+1}$ (where we define $\tau_{n+1} = \tau_{n+2} = \infty$ for convenience). This choice forces the relative order of all peaks with respect to the valleys. There are $\lfloor \frac{n-1}{2} \rfloor$ possibilities for τ in this case.

Similarly, if both sequences are decreasing, that is, $\tau_{[1,4]} = 2413$ or 3412, then τ is determined once we choose the value of τ_1 relative to the peaks of τ , that is, for which index $1 \leq i \leq \lfloor \frac{n}{2} \rfloor$ we have $\tau_{2i} > \tau_1 > \tau_{2i+2}$ (where we define $\tau_{n+1} = \tau_{n+2} = -\infty$ for convenience). There are $\lfloor \frac{n}{2} \rfloor$ possibilities for τ in this case, with just the case $i = 1$ corresponding to $\tau_{[1,4]} = 3412$.

In total, the number of $\tau \in \mathcal{S}_n$ with $|x(\tau)| = n - 2$ is

$$2 \left(2 + \left\lfloor \frac{n-1}{2} \right\rfloor + \left\lfloor \frac{n}{2} \right\rfloor \right) = 2n + 2.$$

□

We do not have formulas for the other entries of Table 6.1.

Problem 6.3. Find a formula for the entries of Table 6.1.

6.2. Asymptotic behavior of the length of the exterior. Even without having exact formulas, we can study the behavior of $|x(\tau)|$ as the length of τ goes to infinity. For given n , we will write \mathbb{P}_n and \mathbb{E}_n to denote the probability and the expectation of events involving $\tau \in \mathcal{S}_n$ chosen uniformly at random. In other words, τ is a random variable with uniform distribution over \mathcal{S}_n . We also use the notation $g(n) = o(f(n))$ or $g(n) \ll f(n)$ to mean that $\lim_{n \rightarrow \infty} g(n)/f(n) = 0$.

The main result in this subsection is a constant asymptotic upper bound on the expected length of the exterior of a permutation. We show that, as the length of the permutation tends to infinity, the limit of this expected value exists and is between $e - 1$ and e . We start with a lemma that bounds the probability that the length of the exterior is large. A similar result appears in [Per13, Lemma 18].

Lemma 6.4. For $1 \leq m \leq n - 1$,

$$\mathbb{P}_n(|x(\tau)| \geq m) \leq \sum_{i=m}^{\lfloor n/2 \rfloor} \frac{1}{i!} + o(n^{-1}),$$

where the summation is defined to be zero when $m > \lfloor n/2 \rfloor$.

Proof. For $1 \leq i \leq n - 1$, let B_i be the event “ τ has a bifix of length i ,” that is, $\tau_{[1,i]} = \tau_{[n-i+1,n]}$. By definition, $|x(\tau)|$ is the largest i such that B_i holds, and

$$\mathbb{P}_n(|x(\tau)| \geq m) = \mathbb{P}_n\left(\bigcup_{i=m}^{n-1} B_i\right) \leq \sum_{i=m}^{n-1} \mathbb{P}_n(B_i).$$

When $i \leq \lfloor n/2 \rfloor$,

$$\mathbb{P}_n(B_i) = \frac{1}{i!},$$

since among all the possible permutations of the last i entries, B_i holds for exactly one of them.

When $i > \lfloor n/2 \rfloor$, we have the upper bound

$$\mathbb{P}_n(B_i) \leq \frac{1}{(n-i)!^{\lfloor \frac{n}{n-i} \rfloor - 1}}. \quad (6.1)$$

This is because if we break up τ into $\lfloor \frac{n}{n-i} \rfloor$ blocks of $n-i$ entries starting at the beginning: $\tau_{[1,n-i]}, \tau_{[n-i+1,2(n-i)]}, \dots$, then, for B_i to hold, the entries in all of these blocks have to be in the same relative order. For $\lfloor n/2 \rfloor < i \leq n - \lfloor \sqrt{n} \rfloor$, we have that $n-i \geq \lfloor \sqrt{n} \rfloor$ and $\lfloor \frac{n}{n-i} \rfloor \geq 2$, and so expression (6.1) gives

$$\mathbb{P}_n(B_i) \leq \frac{1}{\lfloor \sqrt{n} \rfloor!}.$$

For $n - \lfloor \sqrt{n} \rfloor < i \leq n - 2$, we have that $n-i \geq 2$ and $\lfloor \frac{n}{n-i} \rfloor \geq \lfloor \sqrt{n} \rfloor$, and so

$$\mathbb{P}_n(B_i) \leq \frac{1}{2^{\lfloor \sqrt{n} \rfloor - 1}}.$$

When $i = n - 1$, it is easy to see directly that

$$\mathbb{P}_n(B_{n-1}) = \frac{2}{n!},$$

since by Lemma 2.2, B_{n-1} only holds for the two monotone permutations.

Combining the above bounds, we get

$$\mathbb{P}_n(|x(\tau)| \geq m) \leq \sum_{i=m}^{\lfloor n/2 \rfloor} \frac{1}{i!} + \sum_{i=\lfloor n/2 \rfloor + 1}^{n - \lfloor \sqrt{n} \rfloor} \frac{1}{\lfloor \sqrt{n} \rfloor!} + \sum_{i=n - \lfloor \sqrt{n} \rfloor + 1}^{n-2} \frac{1}{2^{\lfloor \sqrt{n} \rfloor - 1}} + \frac{2}{n!}.$$

Clearly, the terms other than the first summation approach 0 faster than n^{-k} for any positive constant k , as n tends to infinity. \square

Theorem 6.5. *The limit of $\mathbb{E}_n(|x(\tau)|)$ as n goes to infinity exists and*

$$e - 1 \leq \lim_{n \rightarrow \infty} \mathbb{E}_n(|x(\tau)|) \leq e.$$

Proof. By definition of expectation,

$$\begin{aligned} \mathbb{E}_n(|x(\tau)|) &= \sum_{i=1}^{\lfloor n/2 \rfloor} i \mathbb{P}_n(|x(\tau)| = i) + \sum_{i=\lfloor n/2 \rfloor + 1}^{n-1} i \mathbb{P}_n(|x(\tau)| = i) \\ &= \sum_{m=1}^{\lfloor n/2 \rfloor} \mathbb{P}_n(m \leq |x(\tau)| \leq \lfloor n/2 \rfloor) + o(1), \end{aligned} \quad (6.2)$$

where we have bounded the second summation from above by

$$n \mathbb{P}_n(|x(\tau)| \geq \lfloor n/2 \rfloor + 1) = o(1),$$

using Lemma 6.4.

To prove that the limit exists, we will show that $\mathbb{E}_n(|x(\tau)|)$ behaves asymptotically similarly to an increasing sequence, and that it is bounded from above. For n_0 with $1 \leq n_0 < n$, consider the bijection

$$\begin{aligned} \mathcal{S}_n &\longrightarrow \mathcal{S}_{n_0} \times \mathcal{S}_{n-n_0} \times \binom{[n]}{n_0} \\ \tau &\mapsto (\sigma, \pi, A) \end{aligned}$$

defined as follows. Let

$$\sigma = \tau_{[1, \lfloor n_0/2 \rfloor] \cup [n - \lceil n_0/2 \rceil + 1, n]}$$

be the permutation obtained by concatenating the first $\lfloor n_0/2 \rfloor$ and the last $\lceil n_0/2 \rceil$ entries of τ , let A be the set of values of these entries before applying the reduction map ρ , and let

$$\pi = \tau_{[\lfloor n_0/2 \rfloor + 1, n - \lceil n_0/2 \rceil]}$$

be the reduction of the remaining entries. If $|x(\sigma)| \leq \lfloor n_0/2 \rfloor$, then clearly $|x(\tau)| \geq |x(\sigma)|$, since the bifix of σ of length $|x(\sigma)|$ is also a bifix of τ . This means that for any given m , the bijection associates to every triple (σ, π, A) with $m \leq |x(\sigma)| \leq \lfloor n_0/2 \rfloor$ a permutation τ with $|x(\tau)| \geq m$. It follows that

$$|\{\sigma \in \mathcal{S}_{n_0} : m \leq |x(\sigma)| \leq \lfloor n_0/2 \rfloor\}|(n - n_0)! \binom{n}{n_0} \leq |\{\tau \in \mathcal{S}_n : |x(\tau)| \geq m\}|.$$

Dividing both sides by $n!$, we get

$$\mathbb{P}_{n_0}(m \leq |x(\tau)| \leq \lfloor n_0/2 \rfloor) \leq \mathbb{P}_n(|x(\tau)| \geq m) = \mathbb{P}_n(m \leq |x(\tau)| \leq \lfloor n/2 \rfloor) + o(n^{-1})$$

by Lemma 6.4. Summing over m , we have

$$\sum_{m=1}^{\lfloor n_0/2 \rfloor} \mathbb{P}_{n_0}(m \leq |x(\tau)| \leq \lfloor n_0/2 \rfloor) \leq \sum_{m=1}^{\lfloor n/2 \rfloor} \mathbb{P}_n(m \leq |x(\tau)| \leq \lfloor n/2 \rfloor) + o(1).$$

Using now Equation (6.2), it follows that

$$\mathbb{E}_{n_0}(|x(\tau)|) \leq \mathbb{E}_n(|x(\tau)|) + o'(1), \quad (6.3)$$

where here $o'(1)$ is an expression that tends to 0 when both n and n_0 go to infinity.

To show that the sequence $\mathbb{E}_n(|x(\tau)|)$ is bounded from above by a constant, we again use Lemma 6.4 to conclude that

$$\begin{aligned} \mathbb{E}_n(|x(\tau)|) &= \sum_{m=1}^{n-1} \mathbb{P}_n(|x(\tau)| \geq m) \leq \sum_{m=1}^{\lfloor n/2 \rfloor} \sum_{i=m}^{\lfloor n/2 \rfloor} \frac{1}{i!} + o(1) \\ &= \sum_{i=1}^{\lfloor n/2 \rfloor} \frac{i}{i!} + o(1) \leq \sum_{i=1}^{\infty} \frac{1}{(i-1)!} + o(1) = e + o(1). \end{aligned}$$

The above inequality implies that $L := \limsup_n \{\mathbb{E}_n(|x(\tau)|)\}$ exists, and that $L \leq e$. Let us show that L is in fact the limit of the sequence. For any $\epsilon > 0$, choose n_0 so that it satisfies the following two conditions. First, take n_0 large enough so that for $n \geq n_0$, we have $\mathbb{E}_{n_0}(|x(\tau)|) \leq \mathbb{E}_n(|x(\tau)|) + \epsilon/2$ (such an n_0 exists by Equation (6.3)). Secondly, pick n_0 so that $\mathbb{E}_{n_0}(|x(\tau)|) > L - \epsilon/2$ (such an n_0 exists by definition of L). Then, for all $n \geq n_0$,

$$\mathbb{E}_n(|x(\tau)|) \geq \mathbb{E}_{n_0}(|x(\tau)|) - \frac{\epsilon}{2} > L - \epsilon.$$

This proves that

$$\lim_{n \rightarrow \infty} \mathbb{E}_n(|x(\tau)|) = L.$$

To prove the lower bound, recall that B_i is the event “ τ has a bifix of length i .” Clearly $\mathbb{P}_n(|x(\tau)| \geq m) \geq \mathbb{P}_n(B_m)$, which equals $\frac{1}{m!}$ when $1 \leq m \leq \lfloor n/2 \rfloor$. Thus,

$$\mathbb{E}_n(|x(\tau)|) = \sum_{m=1}^{n-1} \mathbb{P}_n(|x(\tau)| \geq m) \geq \sum_{m=1}^{\lfloor n/2 \rfloor} \mathbb{P}_n(B_m) = \sum_{m=1}^{\lfloor n/2 \rfloor} \frac{1}{m!}.$$

The limit of the right-hand side as $n \rightarrow \infty$ is $e - 1$. \square

Problem 6.6. *Find the exact value of $\lim_{n \rightarrow \infty} \mathbb{E}_n(|x(\tau)|)$.*

When $n = 10$, the value of $\mathbb{E}_n(|x(\tau)|)$ is approximately 1.909. Some relatively crude computations (for each n with $11 \leq n \leq 100$, we took a random sampling of $7!$ permutations) suggest that the value of the limit is between 1.9 and 1.92.

Problem 6.7. *Find the limiting distribution of $|x(\tau)|$, that is, find*

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(|x(\tau)| = k)$$

for all k .

It was shown by Bóna [Bón11] that when $k = 1$, the sequence $\mathbb{P}_n(|x(\tau)| = 1)$ is decreasing and so its limit exists, and that $0.3640981 \leq \lim_{n \rightarrow \infty} \mathbb{P}_n(|x(\tau)| = 1) \leq 0.3640993$. Even though for fixed k the sequences $\mathbb{P}_n(|x(\tau)| = k)$ are not monotone in general, it is possible to obtain recurrences similar to those in [Bón11] to show that the corresponding limits exist for all k and are positive.

6.3. Permutations with no carrier, and the Möbius function of most intervals. Recall that τ is said to have a carrier element if $x(\tau) \not\leq i(\tau)$. Carrier elements, defined in [BFS11], play a crucial role in determining the Möbius function of intervals in \mathcal{P} , as we pointed out in Subsection 2.2. A question that arises is how many permutations have a carrier element. In this subsection we study this question, and we use our findings to deduce that most intervals (in a precise sense that we will describe) have a zero Möbius function. Since this subsection involves inequalities in both \mathbb{R} and \mathcal{P} , we will use the symbol $\leq_{\mathcal{P}}$ in the latter case as a distinguisher.

For n from 2 to 10, the number of permutations $\tau \in \mathcal{S}_n$ with no carrier element, that is, satisfying $x(\tau) \leq_{\mathcal{P}} i(\tau)$, is given by the sequence

$$0, 4, 12, 84, 548, 4172, 33984, 315800, 3213032, \dots \quad (6.4)$$

Problem 6.8. *Find a formula for the number of permutations with no carrier element.*

We point out that the displayed terms of the sequence (6.4) are all divisible by four, although we have not proved if this is the case for all terms. Note that permutations with $n \geq 3$ and no carrier element include non-overlapping permutations, since an exterior of length one is trivially contained in the interior, and that the number of such permutations is divisible by 4 by Lemma 6.1. From an asymptotic perspective, comparing the sequence (6.4) with $n!$ suggests that, as n grows, most permutations in \mathcal{S}_n have no carrier element. Next we show that this is indeed the case.

Theorem 6.9.

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} i(\tau)) = 0.$$

Proof. Let $n = |\tau|$ and $k = |x(\tau)|$. If $k < n/2$, let $\bar{x}(\tau) = \tau_{[k+1, n-k]}$ be the permutation obtained from τ by deleting the initial and final occurrences of $x(\tau)$. If $k \geq n/2$, define $\bar{x}(\tau)$ to be the empty permutation (and let us add the empty permutation to \mathcal{P} as its new bottom element for convenience). Since $\bar{x}(\tau) \leq_{\mathcal{P}} i(\tau)$, we have

$$\mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} i(\tau)) \leq \mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} \bar{x}(\tau)).$$

For every m , conditioning on whether $|x(\tau)| \leq m$ or not, we can write $\mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} \bar{x}(\tau))$ as a sum of two terms, the first being

$$\begin{aligned} & \mathbb{P}_n(|x(\tau)| \leq m) \mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} \bar{x}(\tau) \mid |x(\tau)| \leq m) \\ & \leq \mathbb{P}_n(|x(\tau)| \leq m) \mathbb{P}_n(\tau_{[1,m]} \not\leq_{\mathcal{P}} \tau_{[m+1,n-m]} \mid |x(\tau)| \leq m) \\ & \leq \mathbb{P}_n(\tau_{[1,m]} \not\leq_{\mathcal{P}} \tau_{[m+1,n-m]}), \end{aligned} \tag{6.5}$$

and the second being

$$\begin{aligned} & \mathbb{P}_n(|x(\tau)| > m) \mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} \bar{x}(\tau) \mid |x(\tau)| > m) \\ & \leq \mathbb{P}_n(|x(\tau)| > m) \\ & \leq \sum_{i=m+1}^{\infty} \frac{1}{i!} + o(n^{-1}), \end{aligned}$$

using Lemma 6.4.

To bound (6.5), notice that $\tau_{[1,m]}$ and $\tau_{[m+1,n-m]}$ are independent random permutations, and split $\tau_{[m+1,n-m]}$ into $\lfloor \frac{n-2m}{m} \rfloor$ disjoint blocks of size m (with possibly some leftover entries that do not belong to any of the blocks). The probability that each individual block is not an occurrence of $\tau_{[1,m]}$ is $1 - 1/m!$. For $\tau_{[1,m]} \not\leq_{\mathcal{P}} \tau_{[m+1,n-m]}$ to hold, none of the blocks can be an occurrence of $\tau_{[1,m]}$. Since these are independent events, we have that

$$\mathbb{P}_n(\tau_{[1,m]} \not\leq_{\mathcal{P}} \tau_{[m+1,n-m]}) \leq \left(1 - \frac{1}{m!}\right)^{\lfloor \frac{n}{m} \rfloor - 2}.$$

Combining all these bounds, we get

$$\mathbb{P}_n(x(\tau) \not\leq_{\mathcal{P}} i(\tau)) \leq \left(1 - \frac{1}{m!}\right)^{\lfloor \frac{n}{m} \rfloor - 2} + \sum_{i=m+1}^{\infty} \frac{1}{i!} + o(n^{-1}) \tag{6.6}$$

for every m .

Finally, it is enough to choose m to be a slowly-growing function of n with $m(n) \rightarrow \infty$ as $n \rightarrow \infty$, such as $m = m(n) = \lfloor \log \log n \rfloor$. For such choice of m , writing

$$\left(1 - \frac{1}{m!}\right)^{\frac{n}{m}} = \left[\left(1 - \frac{1}{m!}\right)^{m!}\right]^{\frac{n}{m m!}},$$

the term inside the square brackets approaches e^{-1} as $n \rightarrow \infty$. Using Stirling's formula, we see that

$$\begin{aligned} (\log \log n)(\log \log n)! &\ll (\log \log n)^2 \frac{(\log \log n)^{\log \log n}}{e^{\log \log n}} \\ &\ll e^{(\log \log \log n)(\log \log n)} \ll e^{\log n} = n, \end{aligned} \quad (6.7)$$

and so the exponent $\frac{n}{m m!}$ goes to infinity as n grows. It follows that

$$\left(1 - \frac{1}{m!}\right)^{\lfloor \frac{n}{m} \rfloor - 2} \leq \left(1 - \frac{1}{m!}\right)^{\frac{n}{m} - 3} \rightarrow 0$$

as $n \rightarrow \infty$. Clearly, the summation of $1/i!$ in (6.6) tends to 0 as $m \rightarrow \infty$. Thus, the whole right-hand side of (6.6) tends to 0 as $n \rightarrow \infty$, proving the statement. \square

A consequence of Theorem 6.9 is that for any fixed σ , most intervals of the form $[\sigma, \tau]$ will have zero Möbius function.

Corollary 6.10. *For fixed σ , let $\mathbb{P}_n^\sigma(\mu(\sigma, \tau) = 0)$ denote the probability that $\mu(\sigma, \tau) = 0$ when τ is chosen uniformly at random among permutations in \mathcal{S}_n that contain σ . Then, for every σ ,*

$$\lim_{n \rightarrow \infty} \mathbb{P}_n^\sigma(\mu(\sigma, \tau) = 0) = 1.$$

Proof. When $|\tau| - |\sigma| > 2$, Theorem 2.1 states that $\mu(\sigma, \tau) = 0$ unless $\sigma \leq_P x(\tau) \not\leq_P i(\tau)$. By Theorem 6.9, the probability that $x(\tau) \not\leq_P i(\tau)$ when τ is chosen uniformly at random from \mathcal{S}_n tends to 0 as $n \rightarrow \infty$. It follows that, for any given σ ,

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(\mu(\sigma, \tau) = 0) = 1.$$

To restrict to permutations τ containing σ , note that

$$\mathbb{P}_n^\sigma(\mu(\sigma, \tau) \neq 0) = \mathbb{P}_n(\mu(\sigma, \tau) \neq 0 \mid \sigma \leq_P \tau) \leq \frac{\mathbb{P}_n(\mu(\sigma, \tau) \neq 0)}{\mathbb{P}_n(\sigma \leq_P \tau)}.$$

By Lemma 3.7, $\lim_{n \rightarrow \infty} \mathbb{P}_n(\sigma \leq_P \tau) = 1$, so $\lim_{n \rightarrow \infty} \mathbb{P}_n^\sigma(\mu(\sigma, \tau) \neq 0) = 0$ as claimed. \square

We point out that Bernini et al. [BFS11] give several conditions on σ and τ that imply $\mu(\sigma, \tau) = 0$. For example, they show that this happens when the first two entries of τ are not involved in any occurrence of σ . While the results in [BFS11] show that the probability that $\mu(\sigma, \tau) = 0$ is large, their conditions depend on σ , and thus they cannot be used to prove that for fixed σ this probability gets arbitrarily close to one as the length of τ increases, as we do in Corollary 6.10.

ACKNOWLEDGEMENT

We thank anonymous referees for suggestions that improved the exposition.

REFERENCES

- [AEK08] José María Amigó, Sergi Elizalde, and Matthew B. Kennel. Forbidden patterns and shift systems. *J. Combin. Theory Ser. A*, 115(3):485–504, 2008.
- [BFS11] Antonio Bernini, Luca Ferrari, and Einar Steingrímsson. The Möbius function of the consecutive pattern poset. *Electron. J. Combin.*, 18(1):Paper 146, 12 pp., 2011.
- [BJJS11] Alexander Burstein, Vít Jelínek, Eva Jelínková, and Einar Steingrímsson. The Möbius function of separable and decomposable permutations. *J. Combin. Theory Ser. A*, 118(8):2346–2364, 2011.

- [Bjö80] Anders Björner. Shellable and Cohen-Macaulay partially ordered sets. *Trans. Amer. Math. Soc.*, 260(1):159–183, 1980.
- [BM99] Louis J. Billera and Amy N. Myers. Shellability of interval orders. *Order*, 15(2):113–117, 1998/99.
- [Bón11] Miklós Bóna. Non-overlapping permutation patterns. *Pure Math. Appl. (PU.M.A.)*, 22(2):99–105, 2011.
- [Bón12] Miklós Bóna. *Combinatorics of permutations*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, second edition, 2012. With a foreword by Richard Stanley.
- [BW82] Anders Björner and Michelle Wachs. Bruhat order of Coxeter groups and shellability. *Adv. in Math.*, 43(1):87–100, 1982.
- [BW83] Anders Björner and Michelle Wachs. On lexicographically shellable posets. *Trans. Amer. Math. Soc.*, 277(1):323–341, 1983.
- [Eli09] Sergi Elizalde. The number of permutations realized by a shift. *SIAM J. Discrete Math.*, 23(2):765–786, 2009.
- [Eli13] Sergi Elizalde. The most and the least avoided consecutive patterns. *Proc. Lond. Math. Soc. (3)*, 106(5):957–979, 2013.
- [Eli16] Sergi Elizalde. A survey of consecutive patterns in permutations. In *Recent trends in combinatorics*, volume 159 of *IMA Vol. Math. Appl.*, pages 601–618. Springer, [Cham], 2016.
- [EN03] Sergi Elizalde and Marc Noy. Consecutive patterns in permutations. *Adv. in Appl. Math.*, 30(1-2):110–125, 2003. Formal power series and algebraic combinatorics (Scottsdale, AZ, 2001).
- [EN12] Sergi Elizalde and Marc Noy. Clusters, generating functions and asymptotics for consecutive patterns in permutations. *Adv. in Appl. Math.*, 49(3-5):351–374, 2012.
- [Gri80] Jerrold R. Griggs. On chains and Sperner k -families in ranked posets. *J. Combin. Theory Ser. A*, 28(2):156–168, 1980.
- [GSS80] Jerrold R. Griggs, Dean Sturtevant, and Michael Saks. On chains and Sperner k -families in ranked posets. II. *J. Combin. Theory Ser. A*, 29(3):391–394, 1980.
- [HS15] Joshua Hallam and Bruce Sagan. Factoring the characteristic polynomial of a lattice. *J. Combin. Theory Ser. A*, 136:39–63, 2015.
- [Kit11] Sergey Kitaev. *Patterns in permutations and words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg, 2011. With a foreword by Jeffrey B. Remmel.
- [Lec97] Bruno Leclerc. Families of chains of a poset and Sperner properties. *Discrete Math.*, 165/166:461–468, 1997. Graphs and combinatorics (Marseille, 1995).
- [MR06] Anthony Mendes and Jeffrey Remmel. Permutations and words counted by consecutive patterns. *Adv. in Appl. Math.*, 37(4):443–480, 2006.
- [MS15] Peter R. W. McNamara and Einar Steingrímsson. On the topology of the permutation pattern poset. *J. Combin. Theory Ser. A*, 134:1–35, 2015.
- [Nak11] Brian Nakamura. Computational approaches to consecutive pattern avoidance in permutations. *Pure Math. Appl. (PU.M.A.)*, 22(2):253–268, 2011.
- [Per13] Guillem Perarnau. A probabilistic approach to consecutive pattern avoiding in permutations. *J. Combin. Theory Ser. A*, 120(5):998–1011, 2013.
- [Smi14] Jason P. Smith. On the Möbius function of permutations with one descent. *Electron. J. Combin.*, 21(2):Paper 2.11, 19, 2014.
- [Smi15] Jason P. Smith. A formula for the Möbius function of the permutation poset based on a topological decomposition. arXiv:1506.04406 [math.CO], 2015.
- [Smi16] Jason P. Smith. Intervals of permutations with a fixed number of descents are shellable. *Discrete Math.*, 339(1):118–126, 2016.
- [ST10] Einar Steingrímsson and Bridget Eileen Tenner. The Möbius function of the permutation pattern poset. *J. Comb.*, 1(1):39–52, 2010.
- [Ste13] Einar Steingrímsson. Some open problems on permutation patterns. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 239–263. Cambridge Univ. Press, Cambridge, 2013.
- [SV06] Bruce E. Sagan and Vincent Vatter. The Möbius function of a composition poset. *J. Algebraic Combin.*, 24(2):117–136, 2006.

- [SW12] Bruce E. Sagan and Robert Willenbring. Discrete Morse theory and the consecutive pattern poset. *J. Algebraic Combin.*, 36(4):501–514, 2012.
- [Wac99] M. L. Wachs. Obstructions to shellability. *Discrete Comput. Geom.*, 22(1):95–103, 1999.
- [Wac07] Michelle L. Wachs. Poset topology: tools and applications. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 497–615. Amer. Math. Soc., Providence, RI, 2007.
- [Wil02] Herbert S. Wilf. The patterns of permutations. *Discrete Math.*, 257(2-3):575–583, 2002.

DEPARTMENT OF MATHEMATICS, DARTMOUTH COLLEGE, HANOVER, NH 03755, USA
E-mail address: sergi.elizalde@dartmouth.edu

DEPARTMENT OF MATHEMATICS, BUCKNELL UNIVERSITY, LEWISBURG, PA 17837, USA
E-mail address: peter.mcnamara@bucknell.edu